

DAA Ames

NCC2-387

14P.

Fairness of Channel Access for Non-Time-Critical Traffic Using the FDDI Token Ring Protocol

IN-16847

Marjory J. Johnson

March 1986

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 86.9

(NASA-TM-89389) FAIRNESS OF CHANNEL ACCESS
FOR NON-TIME-CRITICAL TRAFFIC USING THE FDDI
TOKEN RING PROTOCOL (NASA) 14 p CSCL 12B

N86-29606

Unclas

G3/66 43344

RIACS

Research Institute for Advanced Computer Science

Fairness of Channel Access for Non-Time-Critical Traffic Using the FDDI Token Ring Protocol

Marjory J. Johnson

March, 1986

Research Institute for Advanced Computer Science
NASA Ames Research Center

The Fiber Distributed Data Interface (FDDI) is an ANSI draft proposed standard for a 100 megabit per second fiber-optic token ring. FDDI supports two types of traffic, synchronous and asynchronous. Synchronous traffic is time-critical traffic; stations are assigned guaranteed bandwidth to support their synchronous needs. Asynchronous traffic is lower priority and is sent only if time permits. The purpose of this paper is to prove analytically that the FDDI access protocol provides all stations on the ring with equal access to the channel to transmit asynchronous frames, regardless of the relative sizes of synchronous bandwidth allocations for individual stations. Analytic results are supported with data from simulation runs.

This work was supported by NASA Contract NAS2-11530 and by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration (NASA) to the Universities Space Research Association (USRA).

Fairness of Channel Access for Non-Time-Critical Traffic Using the FDDI Token Ring Protocol

*Marjory J. Johnson
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, California 94035
USA*

1. Introduction

NASA's requirements for the local area network (LAN) on the Space Station include support for diverse types of traffic. A media access protocol for the Space Station LAN must provide low delay and deterministic access to the channel for real-time control functions (e.g., navigation control and environmental control), while at the same time ensuring eventual access to the channel for traffic of a non-time-critical nature. That is, high priority frames must not prevent low priority frames from eventually being transmitted.

A candidate media access control (MAC) protocol for use on the Space Station is the Fiber Distributed Data Interface (FDDI), an ANSI draft proposed standard for a 100 megabit per second fiber-optic token ring. FDDI supports two types of traffic, synchronous and asynchronous. Synchronous traffic is time-critical traffic, such as real-time control and possibly voice transmission, that requires guaranteed access to the channel; all other traffic is called asynchronous.

A notable feature of the FDDI protocol is its ability to guarantee bandwidth for synchronous traffic. Even though asynchronous traffic is non-

time-critical, it is nevertheless worthwhile to determine the quality of service provided for transmission of asynchronous frames (especially, in consideration of the NASA requirement that even low priority traffic must eventually be serviced).

The purpose of this paper is to show that the FDDI access protocol provides all stations on the ring with equal access to the channel to transmit asynchronous frames, regardless of the particular assignment of synchronous bandwidth. We present an analytic proof, given some simplifying assumptions, and then support the result with data from simulation runs.

2. FDDI Access Protocol

The media access control protocol for FDDI is a timed token rotation protocol. Station (i.e., network interface unit) timers cooperatively attempt to maintain a specified token rotation time by using the observed network load to regulate the amount of time that a station may transmit. As part of the ring initialization process, stations negotiate a value for T_Opr , a parameter which specifies the expected token rotation time. Each station is assigned a percentage of T_Opr for its synchronous transmission. Asynchronous traffic is sent only if the load on the ring is light enough to support it.

Each station has two internal timers which control access to the network, a Token Rotation Timer (TRT) and a Token Holding Timer (THT). A station's TRT is initially loaded with T_Opr and it is reset to T_Opr whenever it expires. The token is said to arrive at a station on time if its TRT has not expired since the station last received the token. Otherwise, the token is said to be late. Specific rules governing a station's actions when it receives the token are as follows.

1. If the token is on time, then the current value of TRT is placed in THT,

and TRT is reset to T_Opr . Both asynchronous and synchronous frames may be transmitted. We assume, reasonably, that a station's synchronous transmission precedes its asynchronous transmission.

2. If the token is late, the TRT is not reset. In this case the TRT is said to accumulate its lateness, and only synchronous frames may be transmitted.
3. TRT is enabled during transmission of all frames, synchronous and asynchronous. THT is enabled only during transmission of asynchronous frames.
4. The length of time an individual station may transmit synchronous frames may not exceed its synchronous bandwidth allocation. The length of time that a station may transmit asynchronous frames is governed by its THT; no asynchronous frames may be transmitted after its expiration. No frames of either class may be transmitted after expiration of the station's TRT.
5. If either the THT or the TRT expires while the station is transmitting, it completes transmission of the current frame before issuing the token.

For more information about the protocol see [1,2,3,4,5,6,7,9].

3. Fairness of Access

The FDDI MAC protocol explicitly guarantees access to the communications channel for synchronous traffic. Asynchronous traffic has lower priority. In fact, if all the stations' synchronous bandwidth allocations add up to 100 percent of T_Opr , asynchronous traffic will be virtually blocked from the ring. We assume for the remainder of this paper that synchronous assignments are sufficiently less than 100 percent of T_Opr , so as to allow asynchronous transmission to occur. In addition we assume that all asynchronous frames have the

same priority.

It is not difficult to show that the FDDI protocol provides equal access to the communications channel if all traffic on the network is asynchronous, i.e., if all synchronous bandwidth assignments are zero. The presence of synchronous traffic clearly limits access to the channel for asynchronous transmission.

Surprisingly, traffic patterns are still predictable under some reasonable assumptions about the constancy of synchronous traffic, and access to the communications channel for asynchronous transmission is still fair.

3.1. Ideal Situation, Assuming No Overhead

Throughout this section, we assume zero overhead, i.e., propagation delay and internal station delay are both zero, stations abruptly cease transmission upon expiration of either the THT or the TRT, token transmission time is zero, etc. To make our results independent of the pattern of generation of frames, we assume that an infinite supply of asynchronous frames is available for transmission at each station on the ring. Finally, we assume that each station uses its entire synchronous bandwidth allocation each time it receives the token. Due to the time-critical nature of synchronous transmission, such constancy of synchronous traffic seems to be a reasonable assumption. Let S be the total bandwidth allotted for synchronous traffic from all stations on the ring.

We use the following terminology. A cycle is one complete rotation of the token. We define an *extended cycle* to be a block of $n + 1$ adjacent turns on the ring, where n is the total number of stations. That is, an extended cycle begins and ends with the same station; it is a cycle, extended to include the next station's access to the channel.

Lemma 1: The maximum time available for asynchronous transmission during any extended cycle is $T_Opr - S$.

Proof: According to the FDDI protocol, when a station receives the token on time, the amount of time it may transmit asynchronous frames is equal to the amount of time which remains before its TRT expires. If the token arrived on time the last time around, causing the station's TRT to be reset, this value is $T_Opr - C$, where C is the transmission time during the immediately preceding cycle; otherwise, it is less. Since each station uses its entire synchronous bandwidth allocation during each token rotation, $C \geq S$. Hence, the maximum amount of time left on any station's TRT when it receives the token is $T_Opr - S$. Since any asynchronous transmission which occurs during an extended cycle causes a direct reduction in the amount of time available for asynchronous transmission by downstream stations during the same extended cycle, $T_Opr - S$ is the maximum amount of time available for asynchronous transmission by all stations during a single extended cycle.

Corollary 2: Each cycle has length $\leq T_Opr$. That is, the token will never arrive late at a station.

The next result establishes a relationship between the amount of time a station transmits asynchronous frames during a given turn to the amount of time its upstream neighbor transmitted asynchronous frames during its turn before last. Suppose there are n stations on the ring, labeled $0, 1, \dots, n-1$. To simplify notation we denote station m 's upstream neighbor by station $m-1$, rather than the more correct terminology, $(m-1) \bmod n$. In addition, to simplify the precise statement of the result, we assume that for a given m , station turns are numbered so that station $(m-1)$'s x^{th} turn precedes station m 's x^{th} turn.

Lemma 3: The amount of time that station m , $m = 0, \dots, n-1$, transmits asynchronous frames during its c^{th} turn, $c \neq 1$, is exactly the amount of time that station $(m-1)$ transmitted asynchronous frames during its $(c-1)^{st}$ turn.

Proof: Since the token never arrives late at a station, the amount of time available for transmission of asynchronous frames during each extended cycle is exactly $T_Opr - S$. Since we are assuming an infinite source of asynchronous frames at each station, each station transmits the maximum amount of data allowed. Hence, the length of time used for asynchronous transmission during each extended cycle is exactly $T_Opr - S$. Since the extended cycle starting with station $m-1$'s $c-1^{st}$ turn and the overlapping extended cycle ending with station m 's c^{th} turn differ only by station $m-1$'s $c-1^{st}$ turn on one end and station m 's c^{th} turn on the other, the amount of asynchronous transmission during these two turns must be equal.

Theorem 4: All stations have equal access to the ring to transmit asynchronous frames.

Proof: By Lemma 3, the amount of time an individual station transmits asynchronous frames shifts cyclically around the ring. Hence, after $n+1$ token rotations, where n is the number of stations on the ring, each station has received exactly the same amount of time to transmit asynchronous frames.

It is important to note that the above theorem is independent of the particular assignment of synchronous bandwidth to the separate stations. That is, regardless of the relative sizes of synchronous bandwidth allocations for individual stations, if the ring is observed for a sufficiently long period of time (i.e., for a period of time that is large compared to $n+1$ times the average cycle time, where n is the number of stations on the ring), then each station receives an

equal share of the total time available for asynchronous transmission.

3.2. Realistic Situation, Including Overhead

The two major simplifying assumptions made in the previous section were the abrupt ceasing of transmission upon timer expiration and constancy of synchronous traffic. In this section we discuss the effect of each of these assumptions on the validity of our analytic results.

If a station is transmitting an asynchronous frame when either the THT or the TRT expires, the station completes transmission of that frame before passing the token to the next station. Such a transmission overrun would detract from the predicted amount of time available for transmission of asynchronous frames by downstream stations. Overruns occur with equal probability at all stations. Thus, if asynchronous frame size is constant, all stations should be affected equally by this phenomenon and the validity of our analytic results should not be affected significantly.

A station's inconstant use of its synchronous bandwidth allocation clearly affects the amount of time downstream stations have available for asynchronous transmission, so that the pattern of asynchronous transmission described in Lemma 3 would be only approximate. It is reasonable to expect that a station's use of its synchronous bandwidth allocation will be relatively constant. Hence, minor increases and decreases in the amount of time available for asynchronous transmission caused by inconstant synchronous bandwidth usage should cancel each other's effects in the long run. That is, inconstant use of synchronous allocation should have a minimal effect on our analytic results.

Other overhead factors mentioned in the previous section are minimal and

should have no noticeable affect on our analytic results. See [5] for a detailed discussion of how these factors limit the maximum amount of both asynchronous and synchronous transmission allowed during each token rotation.

At NASA Ames Research Center we have developed a simulation which models the FDDI MAC protocol. Simulation runs support our analytic results. Each table below categorizes channel access during a single simulation run, according to station and transmission type. The first two tables illustrate the fact that fairness of access to transmit asynchronous frames is independent of use of the channel for synchronous transmission. Table 1 categorizes channel access when no synchronous traffic is present on the ring. Table 2 categorizes channel access when individual stations have different synchronous bandwidth allocations; each station uses its entire synchronous bandwidth allocation every time it receives the token.

Percentages presented in these tables measure use of the channel as a percentage of the length of the observation period. Total channel access adds up to less than 100% because of overhead. In both cases, the asynchronous load was virtually infinite at each station. Also, asynchronous frame size was constant, so as to minimize the effects of transmission overrun on channel usage. Careful study of the detailed output from the simulations reveals the pattern of asynchronous transmission described in Lemma 3.

Access to Channel		
Station	Sync Transmission (%)	Async Transmission (%)
1	0	15.2
2	0	14.9
3	0	14.3
4	0	14.6
5	0	14.9
6	0	15.2

Table 1.

Access to Channel		
Station	Sync Transmission (%)	Async Transmission (%)
1	8.8	8.5
2	8.8	8.8
3	0.0	9.3
4	17.9	8.8
5	0.0	8.9
6	0.0	8.6

Table 2.

Additional issues which affect fairness of channel access have been identified via the simulation. In addition to the FDDI MAC protocol, our simulation models some buffer management, in both the transmitter and the receiver. At the transmitter we model a link queue, of length ≥ 0 (to be specified by the user), and at the receiver we model management of buffers to accept frames delivered to the station.

If the link queue is of length > 0 , frames are placed into the queue when they are ready for transmission, and they remain there until they are transmitted on the channel. Since this queue is never purged, frames which are unsendable but which are at the head of the queue may block other frames which could legitimately be transmitted according to the FDDI protocol. This phenomenon is illustrated in Table 3. In this case link queues of length 2 were specified. The

apparent unfairness of access for station 1 is caused by synchronous frames filling up its link queue, thus preventing asynchronous frames from being transmitted even though sufficient time remains on the Token Holding Timer. A similar problem would arise if we modeled multiple priority levels for asynchronous frames, with lower priority asynchronous frames blocking higher priority asynchronous frames.

Access to Channel		
Station	Sync Transmission (%)	Async Transmission (%)
1	17.1	3.4
2	0.0	13.1
3	0.0	13.6
4	0.0	13.6
5	0.0	13.7
6	0.0	13.4

Table 3.

In an attempt to model end-to-end network usage, our simulation models a buffer at the receiving station which stores a frame until the end application (layer 7 of the International Standards Organization's (ISO) Open Systems Interconnection (OSI) network reference model) is able to process it. Frames are processed one at a time. Both the size of this buffer and the amount of time it takes the application to process a frame are user-specified parameters. During the simulation, if there is no room in the buffer to store an incoming frame, it is refused. Simulation results show that buffer limitations in the receiver adversely affect fairness of channel access; stations will have equal access to transmit frames, but will not be able to successfully complete transmission on an equal basis. This situation is similar to that described in [8].

Clearly, fairness of access encompasses more than just fairness of opportunity to access the channel. It also must include fairness to complete the

transmission once it has begun. We intend to investigate these issues more thoroughly in the future.

4. Conclusion

We have proved analytically, given some simplifying assumptions, that the FDDI MAC protocol provides equal access to the communications channel for transmission of asynchronous frames, regardless of the relative sizes of individual stations' synchronous bandwidth allocations. Data from simulation runs supports our analytic results. Finally, we identify other issues which affect the fairness of any media access control protocol, and we illustrate the effect of these issues via simulation runs.

References

1. "FDDI Token Ring Media Access Control," Draft Proposed American National Standard, X3T9.5/83-16, Rev 8, Mar. 1, 1985.
2. R. M. Grow, "A Timed Token Protocol for Local Area Networks," *Electro/82*, Token Access Protocols (17/3), May 1982.
3. V. Iyer and S. P. Joshi, "Streamlining Protocols at High Data Rates," *Proceedings of Wescon/84*, Session 2, Paper 4.
4. V. Iyer and S. P. Joshi, "FDDI's 100 M-bps Protocol Improves on 802.5 Spec's 4-M-bps Limit," *EDN*, pp. 151-160, May 2, 1985.
5. M. J. Johnson, "Proof that Timing Requirements of the FDDI Token Ring Protocol are Satisfied," *RIACS Technical Report 85.8*, August, 1985.
6. S. P. Joshi and V. Iyer, "New Standards for Local Networks Push Upper Limits for Lightwave Data," *Data Communications*, pp. 127-138, Jul. 1984.
7. F. E. Ross, and R. K. Moulton, "FDDI Overview - A 100 Megabit per Second Solution," *Proceedings of Wescon/84*, Session 2, Paper 1.
8. J. Sventek, W. Greiman, M. O'Dell, and A. Jansen, "Token Ring Local Area Networks: A Comparison of Experimental and Theoretical Performance," *Computer Networks*, Vol. 8, No. 4, Aug. 1984, pp. 301-309.
9. J. M. Ulm, "A Timed Token Ring Local Area Network and its Performance Characteristics," *IEEE*, Feb. 1982.



Mail Stop 230-5
NASA Ames Research Center
Moffett Field, CA 94035
(415) 694-6363

The Research Institute for Advanced Computer Science
is operated by
Universities Space Research Association
The American City Building
Suite 311
Columbia, MD 21044
(301) 730-2656